

Time-domain simulation of nonlinear radiofrequency phenomena

Thomas G. Jenkins,¹ Travis M. Austin,¹ David N. Smithe,¹ John Loverich,¹ and Ammar H. Hakim² ¹*Tech-X Corporation, 5621 Arapahoe Avenue, Boulder, Colorado 80303, USA*

²Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543, USA

(Received 17 October 2012; accepted 2 January 2013; published online 17 January 2013)

Nonlinear effects associated with the physics of radiofrequency wave propagation through a plasma are investigated numerically in the time domain, using both fluid and particle-in-cell (PIC) methods. We find favorable comparisons between parametric decay instability scenarios observed on the Alcator C-MOD experiment [J. C. Rost, M. Porkolab, and R. L. Boivin, Phys. Plasmas **9**, 1262 (2002)] and PIC models. The capability of fluid models to capture important nonlinear effects characteristic of wave-plasma interaction (frequency doubling, cyclotron resonant absorption) is also demonstrated. © 2013 American Institute of Physics. [http://dx.doi.org/10.1063/1.4776704]

I. INTRODUCTION

National and international class magnetic fusion energy experiments, including DIII-D, C-Mod, NSTX, and ITER, all rely on radiofrequency (RF) heating as a principal means of achieving the requisite high temperature plasma. In these experiments, the RF power must first pass through the lower density, lower temperature edge region before reaching the high-density core plasma where it is to be deposited. However, non-linear processes are enabled by the relatively high fraction of power density associated with the RF in the edge region. Such nonlinearities may parasitically draw energy from the RF wave, leading to degraded or even interrupted core plasma performance. Non-linear processes of particular concern include the formation of RF sheaths near antenna structures,¹ the generation of higher harmonic waves at multiples of the injected RF wavenumber or frequency, and the development of parametric decay instabilities (PDIs);² the latter two instabilities will be numerically modeled in this work. The PDI is a three-wave coupling process wherein injected RF waves decay into separate daughter waves of different wavelengths and frequencies, in accordance with certain selection rules which we will later discuss. In such cases, the nature of the daughter waves often causes RF power to be deposited primarily in the plasma edge, rather than in the core. PDIs have been observed and analyzed in lower-hybrid,^{3–5} electron-cyclotron,^{6,7} and ion-cyclotron^{8,9} regimes. Their presence has been noted on several major experimental devices, including DIII-D,¹⁰ JT-60U,¹¹ C-Mod,¹² and ASDEX.¹³ Numerical simulations of the PDI in magnetic fusion-relevant regimes include studies of upper hybrid/electron cyclotron resonance heating scenarios,14 the decay of electron Bernstein waves,¹⁵ and the excitation of ion Bernstein waves¹⁶ using particle-in-cell (PIC) codes. In addition, full-wave codes such as AORSA¹⁷ have also been used to model PDI scenarios in the frequency domain.¹⁸

In this work, we will summarize recent developments in time-domain simulation methods for RF physics, both in field models and particle-in-cell models, which have improved quantitative modeling capabilities of non-linear RF phenomena. More experimentally relevant simulations, including two- and three-dimensional geometries and full-spectrum analysis, are made possible by these new methods. Such improved capability will allow for a better understanding of existing experimental observations as well as enabling greater confidence in predictive modeling for future experiments such as ITER. Ultimately, we anticipate that further developments of the techniques presented herein will enable quantitative, three-dimensional, non-linear analysis of RF physics effects in the plasma edge region.

In Sec. II of this work, we summarize and discuss the relevant non-linear physics associated with the PDI. We emphasize the word non-linear, in part because so much of RF theory (especially that associated with megawatt heating on large tokamak experiments) has historically been carried out using linear analysis. Indeed, in Sec. III, we present the generalization of one such linear technique, based upon particle methods, to the non-linear realm. However, we also demonstrate (in Sec. IV) that a non-linear approach based on fluid modeling-an approach less common to RF problems-can be applied to investigate PDI scenarios. Section V discusses results from PDI simulations which use the PIC model and compares these results with experimental data from the Alcator C-MOD device, demonstrating significant agreement. In Sec. VI, we discuss the results of comparable simulations which use the fluid model and demonstrate the non-linear generation of higher-harmonic phenomena in the fluid code. Conclusions and future lines of research are presented in Sec. VII.

II. NON-LINEAR EFFECTS AND THE PARAMETRIC DECAY CHANNEL

One measure of non-linearity is to compare the energy density of a wave, $1/2\varepsilon_{plasma}|\mathbf{E}|^2$ (where **E** is the electric field vector and ε_{plasma} is the linear dielectric associated with the wave), to the thermal energy of the background milieu, e.g., nT (where *n* is the density and *T* is the temperature). In the core of a magnetic fusion plasma, one finds that even megawatt RF power levels constitute a small fraction, about 10^{-6} , of the thermal stored energy. Of course, this must be so, or else confinement would be too adversely affected by such power for it to be a practical means of heating plasma.

However, in the edge region, where densities and temperatures are factors of 10^2 to 10^3 lower than in the core, the wave energy can begin to approach the thermal energy, and the assumption of linear behavior can break down.

When wave energy is still relatively low, the primary non-linear interaction is through products of what are still essentially linear waves. (This is the regime in which our later simulations will be performed.) A linear wave has a dispersion relation, which relates the wave frequency ω_0 and wavenumber \mathbf{k}_0 ; other waves, which satisfy their own dispersion relations with frequency/wavenumber combinations (ω_1 , \mathbf{k}_1) and (ω_2 , \mathbf{k}_2), may also exist in the plasma at the same time. If by chance, these waves satisfy the selection rules

$$\mathbf{k}_0 = \mathbf{k}_1 + \mathbf{k}_2 \text{ and } \omega_0 = \omega_1 + \omega_2, \tag{1}$$

then the three waves can interact and exchange energy. Though it might seem unlikely that all four of the above equations could be satisfied simultaneously, it is in fact quite easy due to the particular nature of many dispersion relations. Many plasma waves are inherently perpendicular, parallel, electrostatic, or transverse, and thus constrain the dimensionality of the wavenumber equation to 1 or 2 dimensions, rather than 3. Furthermore, many waves have resonances and cutoffs, in which small changes in ω result in large changes in k; such behaviors manifest themselves as sections of nearly vertical behavior on a dispersion relation plot, $\mathbf{k}(\omega)$. Consequently, if one can match the frequency equation with such a wave near a resonance or cutoff, one need merely slide up or down this vertical section to find the k which makes the three-wave coupling work. The ion Bernstein wave (IBW) is one such wave that exhibits such near vertical behavior; the ion cyclotron wave (ICW) has similar characteristics. In our simulations, we focus on a fast wave for (ω_0, \mathbf{k}_0) (which we refer to hereafter as the pump wave) and IBW and/or ICW for the daughter waves (ω_1, \mathbf{k}_1) and (ω_2, \mathbf{k}_2) . In Figure 1, we show the IBW dispersion relation for various values of the ratio $R = \omega_{pi}^2 / \Omega_i^2$ (where ω_{pi} and $\Omega_{\rm I}$ are the conventional ion plasma frequency and ion cyclotron frequency), together with red and blue parallelograms



FIG. 1. Illustration of the parametric decay process. The summation of daughter waves matches the pump wave, permitting energy exchange among the three waves. The summation of daughter waves can be illustrated as a parallelogram on the overlaid dispersion relations. Here, R is the square of the ratio of ion plasma frequency to ion gyrofrequency (see also p. 301 of Ref. 19).

demonstrating how ω and **k** values can sum to form daughter pairs for a particular value of R. The Alfvénic dispersion of the fast pump wave is represented as the straight line forming the diagonal of either of the parallelograms (with the differing slopes corresponding to different plasma temperatures). Note how the vertical behavior of the Bernstein waves facilitates the process of matching wavenumbers.

Looking at the blue parallelogram from the figure, we see that in this case, the pump wave is split into two daughters of nearly half frequency and wavenumber. More common is the case of the red parallelogram, where one daughter has a nearzero wavenumber while the other has nearly the same wavenumber as the pump wave. In this case, the daughter waves lie on sections of nearly vertical dispersion behavior, either near resonance or near cutoff. The term "quasi-mode" is often used to describe such waves, as their spatial localization properties are such that they appear only in specific locations or appear with little or no spatial oscillation.⁹

In an experimental device, density gradients in the plasma edge will give rise to variation in the R parameter of Figure 1; a pump wave of fixed frequency passing through the plasma edge must then also experience variation in k_{\perp} if the dispersion relation is to be satisfied. Similarly, temperature variation in the plasma edge will effectively rescale the y-axis in the figure, yielding variation in the allowable k_{\perp} values even when R is fixed. Pump waves at a fixed frequency may thus be susceptible to multiple parametric decay channels as they pass through the edge region. Although Figure 1 does not attempt to capture this complicated path through the parameter space, it nevertheless encapsulates the essential physics which makes the PDI a viable decay channel—namely, the near-vertical behavior (steep k-variation with respect to frequency) of solutions to the dispersion relation at large k_{\perp} .

An important aspect of our simulation techniques is that solutions are calculated in the time domain, rather than in the frequency domain. Though use of the time domain is necessary for very highly non-linear behaviors, in weaker coupling situations (such as the aforementioned three-wave coupling), one can still compute solutions in the frequency domain. In the latter case, pump and daughter waves are selected *a priori* and then their mixing is computed from the wave products.²⁰ However, the frequency domain technique may fail if there are many simultaneous decay products for the same pump wave (as is commonly the case when high harmonics, together with spatially varying magnetic fields, densities, and temperatures, are involved in the wave physics). Even if there is one dominant decay pair, this pair must be identified *a priori*.

We use the time-domain approach in this project, even though the expected degree of non-linearity is still mild. Numerical approaches for both fluid and PIC models exist which can be applied to the problem in the time domain; in this work, we use the recently developed Nautilus fluid code^{21,22} as well as the more mature Vorpal electromagnetic PIC code.²³ In a time-domain approach, all possible decay products of the pump wave can exist simultaneously, and dominance of one over the other is arrived at naturally through application of the fundamental equations, rather than through *a priori* selection.

III. PARTICLE-IN-CELL METHODS

PIC methods²⁴ are ubiquitous in present-day numerical plasma modeling. Such methods find use in the study of plasma wakefield accelerators,²⁵ astrophysical problems,²⁶ and propulsion,²⁷ to name only a few applications. In magnetic fusion research, they have been used to investigate the behavior of ion temperature gradient-driven turbulence²⁸ and the kinetic stabilization of internal kink modes²⁹ in tokamaks, as well as to model wave-particle interactions in the plasma core.³⁰

In the most rudimentary form (the "total-f" method) of PIC, one writes the collisionless kinetic equation for the distribution function f_s

$$\frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \frac{\partial f_s}{\partial \mathbf{x}} + \frac{q_s}{m_s} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f_s}{\partial \mathbf{v}} = 0$$
(2)

and posits that this distribution function can be represented as the statistical average of a Klimontovich distribution

$$f_s(\mathbf{x}, \mathbf{v}, t) = \sum_{i=1}^N \delta[\mathbf{x} - \mathbf{x}_i(t)] \delta[\mathbf{v} - \mathbf{v}_i(t)]$$
(3)

comprised of N computational particles (also designated as "markers" or "Lagrangian markers" in the literature³¹). These objects evolve along characteristic trajectories of the kinetic equation

$$\frac{d\mathbf{x}_i(t)}{dt} = \mathbf{v}_i(t); \quad \frac{d\mathbf{v}_i(t)}{dt} = \frac{q_s}{m_s} (\mathbf{E} + \mathbf{v}_i(t) \times \mathbf{B})$$
(4)

and their initial loading mimics the physical distribution function. If these computational particles comprise a discrete realization of the physical distribution function at the simulation outset, their collective behavior will (as they evolve forward in time along characteristic trajectories) continue to statistically represent this function in the limit of large particle count N.

In many physical scenarios, it is desirable to calculate a small distribution function perturbation f_{1s} about some equilibrium distribution function f_{0s} whose spatial and temporal variations are not relevant to the problem at hand (i.e., it evolves on slower timescales and on broader spatial scales than the perturbations of interest). In such cases, the more general "delta-f" technique may be used; writing the kinetic equation for $f_s = f_{0s} + f_{1s}$ in the form

$$\frac{\partial f_{1s}}{\partial t} + \mathbf{v} \cdot \frac{\partial f_{1s}}{\partial \mathbf{x}} + \frac{q_s}{m_s} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f_{1s}}{\partial \mathbf{v}}$$
$$= -\frac{q_s}{m_s} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f_{0s}}{\partial \mathbf{v}}$$
(5)

enables the perturbed distribution function f_{1s} to likewise be represented as the statistical average of a Klimontovich distribution

$$f_{1s}(\mathbf{x}, \mathbf{v}, t) = \sum_{i=1}^{N} w_i(t) \delta[\mathbf{x} - \mathbf{x}_i(t)] \delta[\mathbf{v} - \mathbf{v}_i(t)].$$
(6)

In addition to the characteristic Vlasov trajectories, one also has an equation describing the evolution of particle weight

$$\frac{d\mathbf{x}_{i}(t)}{dt} = \mathbf{v}_{i}(t); \quad \frac{d\mathbf{v}_{i}(t)}{dt} = \frac{q_{s}}{m_{s}}(\mathbf{E} + \mathbf{v}_{i}(t) \times \mathbf{B});$$
$$\frac{dw_{i}(t)}{dt} = -\frac{1}{g_{i}} \left[\frac{q_{s}}{m_{s}}(\mathbf{E} + \mathbf{v}_{i}(t) \times \mathbf{B}) \cdot \frac{\partial f_{0s}}{\partial \mathbf{v}} \right], \tag{7}$$

where the g_i are discrete realizations of a normalized probability distribution function describing the initial distribution of the computational markers. (This distribution need not be Maxwellian or correspond to the physical distribution function; see Ref. 32). The collective behavior of the computational particles, provided the initial loading is correct, is again an accurate representation of the perturbed distribution function when N is large. One may also utilize the delta-f technique to carry out total-f simulations with arbitrary loading strategies, provided that the initial particle weights (which in this case do not evolve in time but do not necessarily have value unity) are correctly calculated.

When deviations from thermal equilibrium are small, the delta-f technique is essentially a linear technique, but delta-f methods can also be extended into the nonlinear regime to model large deviations from equilibrium if sufficient measures³¹ to counter statistical noise are employed. Because statistical fluctuations scale as $N^{-1/2}$, the most common such measure is to increase the number of computational particles used in the simulation; however, judicious choice of the initial distribution g also can prove advantageous (both in delta-f and total-f approaches). In particular, we have made use in this work of an extremely low-noise particle loading approach inspired by the work of Zu and Qin,³³ wherein particles are loaded uniformly in gyro-angle on a "Kruskal ring"-that is, a ring of particles all having the same parallel and perpendicular velocity, and the same gyrocenter. In the linear regime, invariance properties of the Kruskal ring ensure that the noise-quenching uniform gyroangle spacing is preserved forever, assuring low noise late in the simulation after many RF cycles. An illustration of a failed IBW 1-D benchmark test without this loading, and the same successful test with this loading scheme, is shown in Figure 2. In these simulations a time-centered locally implicit method was used for the electrons,³⁴ thus avoiding the need to resolve electron timescales.

In addition to verifying the predicted invariance and low-noise properties of the Kruskal ring approach in the linear regime, we have also verified that delta-f simulations using this technique can successfully be applied to study non-linear RF problems. In these non-linear cases, the gyroangle uniformity is no longer maintained forever. However, its departure from non-uniformity is observed to be slow, so that noise is still well contained even after 150 RF oscillations at the relevant 2×10^4 V/m field strength.

The principle advantage of the particle approach is that it provides full physical content in terms of all harmonics (all-orders) of IBW, ICW, and other waves. It is also easily run in one, two, or three dimensions with virtually no change in method. While particle approaches are often inherently more CPU-intensive than the fluid models we will discuss in

Downloaded 04 Apr 2013 to 198.125.228.209. This article is copyrighted as indicated in the abstract. Reuse of AIP content is subject to the terms at: http://pop.aip.org/about/rights_and_permissions



FIG. 2. Illustration of the beneficial effect of Kruskal ring low-noise particle loading on an IBW simulation. The Gaussian velocity distribution of particles (top left) cannot sustain the IBW due to particle noise; no coherent IBW-wavelength fluctuations arise in the ensuing electric field (top right). Invariance properties of the Kruskal ring distribution of particles (bottom left) in the linear regime enable coherent oscillations at the predicted IBW wavelength to arise in the electric field (bottom right).

Sec. IV, these approaches typically scale very well on largescale supercomputer parallel architectures. The principle disadvantage of the particle approach is the need for a large number of particles; large parallel processing capability is thus required for experimentally relevant simulations. While 1D particle simulation is straightforward, moving to higher dimensions may lead to excessive computational resource requirements or yield untenable time-to-solution results. Methods making use of modern graphics processing unit (GPU) hardware may increase the speed of such calculations³⁵ as this hardware becomes increasingly available on high-performance computation platforms.

It should also be noted that the PIC methods we have described above do not include collisional effects, which may occur among the plasma species as well as between plasma ions/electrons and neutral particles or vessel walls. Dissipative phenomena such as charge exchange, neutral ionization and recycling, recombination, and sheath physics, which determine the shape of density and temperature profiles in the plasma edge, are not considered in detail in this work, as it is not our purpose to construct detailed models of the edge physics. We note, however, that Coulomb collisions and neutral interactions can be incorporated into PIC models through the use of Monte Carlo techniques.³⁶

IV. FLUID METHODS CAPABLE OF MODELING HIGH HARMONICS

The most effective means of non-linear analysis has often been a fluid approach; a cold fluid method can easily simulate entire tokamak domains. However, the requirement of suitable closures for the fluid equations (a hierarchical set of equations obtained via various velocity moments over the kinetic equation) has prevented their use in the most common RF scenarios that require ion Bernstein waves. In particular, the all-orders linear dielectric approach used in the RF community to look at higher-harmonic IBW's (e.g., in the NSTX device, see Ref. 37) is essentially equivalent to a linear version of what might be called an "all-moments" fluid approach. Each pair of fluid velocity moments retained before truncation and closure contributes one additional cyclotron harmonic resonance to the linear dynamics; retention of the scalar charge and vector current density captures the first harmonic resonance in the cold plasma approximation, while retention of the (second-rank) pressure tensor and (third-rank) heat flux tensor also capture the second harmonic resonance.⁴⁴ Hence, the retention of all-orders is equivalent to the retention of all-moments; when (as in NSTX) there are 18+ harmonics contributing to the dynamics, a moment approach appears wholly untenable. So at first glance, it appears as if fluid methods cannot be useful in looking at parametric decay involving multiple cyclotron harmonics.

A closer examination affords some hope, though, that a fluid approach might be applicable. For example, one common fluid closure is known as the "adiabatic" approximation, and this simply sets to zero the rank-3 velocity moment tensor (often denoted with the letter Q_{ijk} , and called the heat flux tensor), together with all higher-order velocity moment tensors. The result is a system involving mass density ρ_m , velocity **V**, and full pressure tensor \mathbb{P} , consisting of temporal update equations for each, as shown below (in the sourceless, collisionless approximation):

$$\partial_t \rho_m + \nabla \cdot (\rho_m \mathbf{V}) = 0 \quad (continuity),$$
 (8)

$$\partial_t(\rho_m \mathbf{V}) + \nabla \cdot [\rho_m \mathbf{V} \mathbf{V} + \mathbb{P}] - \alpha \rho_m (\mathbf{E} + \mathbf{V} \times \mathbf{B}) = 0$$

(force/momentum equation), (9)

$$\{\partial_t + \mathbf{V} \cdot \nabla\} \mathbb{P} + \mathbb{P}\nabla \cdot \mathbf{V} + [\mathbb{P} \cdot \nabla \mathbf{V}] + [\mathbb{P} \cdot \nabla \mathbf{V}]^t - \alpha \mathbb{P} \times \mathbf{B} + \alpha \mathbf{B} \times \mathbb{P} = -\nabla \cdot \mathbb{Q} = 0 \quad (a diabatic approx.).$$
(10)

Charge to mass ratio is $\alpha = (q/m)$. With the adiabatic approximation, $\nabla \cdot \mathbf{Q} = 0$, and taken together with Maxwell's equations, this is a closed system containing all linear dispersion properties associated with the traditional first-order finite-Larmor approximation commonly encountered in RF analysis.³⁸ In particular, the pressure tensor provides the IBW at the fundamental and 2nd harmonics, but no harmonics above that. Thus, while this approach cannot treat the higher harmonic scenarios of the NSTX tokamak, it can immediately be applied to the classic minority-fundamental and 2nd harmonic heating scenarios³⁹ of relevance to ITER. Because non-linear effects are included in this fluid model, and because the calculations can be carried out in the time domain, it is also capable of treating parametric decay scenarios in certain cases.

V. PARTICLE-IN-CELL PARAMETRIC DECAY SIMULATIONS

As we have noted in the introduction, the PDI can degrade performance of ion cyclotron resonance heating in the vicinity of the launcher antenna structure. We consider the case wherein the launched RF wave breaks down into two daughter waves, usually with one daughter mode frequency near the local cyclotron frequency [generally an ICW] and the other daughter wave being an IBW with wavenumber and frequency set by the necessary frequency and wavenumber selection rules. The process is facilitated by the nearvertical dispersion curve of the ICW, allowing it to take on whatever wavenumber is required to couple with the IBW dispersion relation; this wave is also called the quasi-mode.

In its most general form, the behavior of the PDI cannot be modeled with the linear delta-f PIC method or cold fluid methods; it is inherently a nonlinear process. Furthermore, it is a process that involves sub-harmonic frequencies, rather than simple frequency doubling and the subsequent harmonic cascade that occurs in many other nonlinear problems. Thus, observation of sub-harmonic frequencies in the simulations is a potential signature of the PDI. For the simulations presented here, the use of a full-f PIC model for the ions was required to capture these inherent nonlinearities.

PDIs have been observed in Alcator C-Mod and in NSTX. In the case of the Alcator C-Mod observations, the observed modes involve fundamental and 2nd harmonic cyclotron physics, and the properties of the ensuing IBW daughter wave suggest that it is amenable to simulation by both particle and fluid methods (as discussed in Secs. III and IV) in the nonlinear regime. Thus, we have selected a proof-of-principle scenario that matches the PDIs observed with an RF probe by Rost *et al.*^{12,40} Plasma parameters, including density, magnetic field, and ion temperature, were chosen to match the analysis done in these references; the experimental data are shown in Figure 3.

When comparing time-domain simulation results to observations made in Alcator C-Mod, an important aspect to consider is that while the experimental drive frequency is known, the daughter frequencies (and/or spectrum of frequencies) are not *a priori* fixed. Thus, much like the experimental measurements, our simulations must take a long time-history, FFT the results, and look for signals with frequency other than the drive signal. We will demonstrate that in one such simulation, daughter waves in a C-Mod simulation scenario were observed which strongly resemble measurement data.



FIG. 3. An MIT experiment⁴¹ observed likely PDI behavior in vicinity of the RF antenna (left), as shown by the measured spectrum (right). Two daughter modes are identified, one at the ion cyclotron frequency and the other a propagating ion Bernstein wave. [Reprinted with permission from J. C. Rost, R. L. Boivin, M. Porkolab, J. C. Reardon, and Y. Takase, AIP Conf. Proc. **403**, 85 (1997). Copyright 1997 American Institute of Physics.]

Downloaded 04 Apr 2013 to 198.125.228.209. This article is copyrighted as indicated in the abstract. Reuse of AIP content is subject to the terms at: http://pop.aip.org/about/rights_and_permissions



FIG. 4. Snapshot of the electric field components for the 2D particle PDI simulation. The wave is driven with a J_y excitation on the right, with partially absorptive boundaries on the left and right, and periodic boundaries in the y-direction. Fourier transforms are taken at the location indicated by the black dot.

Our simulations span a small 2D region just outside the antenna given by the $R_{major} \times Z$ coordinate range $[0.903 \,m - 0.907 \,m] \times [-1.67 \times 10^{-4} \,m - 1.67 \times 10^{-4} \,m],$ which corresponds to a thin horizontal slice of the region within the white box of Figure 3. Electrons and deuterium ions have density $n_0\!=\!1.0\times10^{19}\,m^{-3}$ and are immersed in a magnetic field with B = 3.504 T pointing in the direction perpendicular to the R-Z plane. The deuterium temperature is 5 eV, and the electrons are treated as a cold fluid. The value of k_z arises naturally from the thinness of the Z interval (which has only 5 cells) and the tendency for there to be one or two wavelengths in that direction. A single frequency (80 MHz) is driven; 100 particles per cell are used with 208 cells in the main simulation dimension. The simulation proceeds for 10 million time steps, taking approximately 2 days on a 32 node parallel computer. This long simulation length was necessitated by the need to simulate many wave periods (giving sufficient time for the PDI to excite the daughter waves to measurable amplitude) and by the high resolution requirements required to capture the short wavelengths of the daughter waves.

In Figure 4, we show the electric field components observed at a fixed point in time late in the aforementioned simulation. The wave was excited on the right side of the computational domain; the simulation is periodic in the y-direction (corresponding to Z). In the x-direction (corresponding to R), the boundary conditions on the fields correspond to a perfectly conducting wall, while particles which impinge these boundaries are absorbed and removed from the simulation. (These conditions—to which the observed PDI physics does not appear to be highly sensitive—effectively prescribe a small net flux out of the simulation domain.) Each component shows markedly different behavior, which is a strong indication of the different nature of the daughter waves.

Because these simulations are two-dimensional, spatial dependency in the toroidal direction [normal to the plane of Figure 4] was not considered; effectively, the current source in the simulation most resembles the monopole ($N_{\varphi} = 0$) linear mode. Numerically, this configuration readily gives rise to IBWs. However, one should note that an experimental monopole antenna has significant spectral width beyond the $N_{\varphi} = 0$ mode, and due to nonlinearity and mode mixing, linear mode characterization is not really sufficient to understand

the full spectral effects. The experiments which generated the data of Figure 3(b) used a dipole phasing,⁴² with the ensuing toroidal mode spectrum peaking at a linear mode of $N_{\varphi} = 10$. Marked differences in RF fields and potentials obviously arise from different phasing scenarios,⁴³ but the relatively small size of our simulation domain means that we are effectively calculating a local approximation to the field spectra, rather than a particular toroidal mode, in much the same manner as the experimental probes measure local conditions rather than toroidally averaged fields. The inclusion of finite parallel wavenumber would enable the study of full spectrum and phasing effects, but fully capturing the nonlinear PDI behavior would require 3D PIC simulations at significantly higher computational cost. We do not consider such effects in this work, but instead focus on the local nonlinear physics of the PDI within the 2D region.

A time-sequence of the E_x electric field was taken at the location indicated by the arrow on the E_x plot of Figure 4, and the FFT spectrum of this data was then calculated (Figure 5). In the data, two sub-harmonic peaks are observed below the drive frequency, and these peaks occur in close proximity to the peaks observed in the C-Mod measurements (overlaid).⁴⁰ The spectral widths of the simulated and experimentally observed peaks are remarkably similar. An interesting discrepancy also arises in that the ICW quasi-mode frequency appears to be slightly down-shifted in the simulation, while the IBW frequency is correspondingly up-shifted. Another notable observation is the presence of a strong



FIG. 5. Spectrum from the normalized E_x (blue) and E_y (green) components of the PIC simulation overlaid on spectrum from probe measurement (compare with the right panel of Figure 3 and the two leftmost panels of Figure 6). The parametric decay of the initial wave into two daughter waves occurs in both experiment and simulation. [Reprinted with permission from J. C. Rost, R. L. Boivin, M. Porkolab, J. C. Reardon, and Y. Takase, AIP Conf. Proc. **403**, 85 (1997). Copyright 1997, American Institute of Physics.]

zero-frequency signal. It is believed that this zero-frequency signal is an expression of induced flow arising from the PDI interaction. The excitation of such flows may have additional implications for plasma confinement.

Further confirmation of the PDI physical interpretation is gleaned from a detailed look at each of the field components and the relative strengths of the frequency peaks in each component. We note that the driven fast wave at 80 MHz is expected to have E_x and E_y signals of similar magnitude, while the IBW near 50 MHz is an electrostatic wave that should have a very dominant E_x signal. In fact, we see precisely this behavior, as shown in Figure 6 below.

To verify that the spectra observed in the simulation were not unduly affected by perturbations to the non-fluctuating distribution function (because these are total-f simulations), relevant timescales of the problem were compared. The saturation of the E_x amplitude, which occurs on the longest nonlinear timescale directly associated with the PDI, occurs after roughly 5×10^{-7} s (about 40 wave periods of the pump wave, and about half the time interval over which the simulation takes place). Timescale estimates for quasilinear modification of the background distribution function are on the order of milliseconds for these simulation parameters and are thus assumed negligible, and the simulation timestep (of order 10^{-13} s) is clearly sufficient to resolve the IBW, ICW, and pump wave dynamics.

Some consideration was given to the slight frequency discrepancies between the simulation and measured signals. Two hypotheses have been suggested as explanations for this effect. The first is that we note the simulation daughter waves lie closer to the 1/3 and 2/3 sub-harmonic frequencies of drive frequency. One possible reason for this might be that such exact sub-harmonics are artificially favored due to particle noise created by the tent-function particle shape (which arises from bilinear interpolation of the particle quantities onto the grids used to calculate field quantities). The tent-function particle shape contains a slope discontinuity and thus might have artificially high emission at the n = 2and higher harmonics. A second hypothesis is that the observed shift might be due to the use of cold fluid electrons in the simulation. Indeed, in Rost's analysis of the PDI scenario,⁴⁰ it is noted that the location of the ICW peak is expected to shift downward with electron temperature, and the predicted frequency at zero-temperature more closely resembles the frequency observed in the simulations.

We have also considered possible reasons for the discrepancies between the saturated amplitudes of the daughter waves in the experiment and simulation data of Figure 5. Uncertainty in the angular orientation of the probe makes the analysis somewhat challenging; while the probe is aligned so as to be orthogonal to the strong guide field (z-direction), its orientation in the x-y plane is unknown. In our simulations, IBWs are assumed to propagate in the x-direction, and in this coordinate basis, the probe appears to be predominantly measuring what our simulations call E_v (note the similarity between the green curve and the experimental data for frequencies above 20 MHz in Figure 5). However, the lowfrequency components of the experimental spectrum are more characteristic of the larger-amplitude E_x component, so the probe is most likely detecting predominantly E_v , together with a small projection of E_x which produces the low-frequency components of the spectrum. While the comparable power content in the two daughter waves of the E_y spectrum is slightly larger ($\sim 10-15 \, dB$) than that of the daughter waves observed in the experiment (when the simulated and experimental power content of the drive wave are equal), this is a nonlinear effect; the ratio of daughter-todrive power content is certainly negligible for low drive amplitude (when the PDI does not occur) but is raised nonlinearly as this amplitude is increased. Thus, we conclude that the RF power in our simulations was probably slightly higher than was present in the C-MOD experiment, and that the experimental probe was likely very nearly orthogonal to the direction of the propagating IBW. The discrepancy between power levels may be related to an observation by C-MOD experimentalists that the amplitude of spectral peaks is sensitive to the probe positioning relative to the poloidal limiters (see Section 6.4.1 of Ref. 42), an effect which is not captured by our two-dimensional simulation geometry. It may also arise from some weakly dissipative mechanism (neutral or Coulomb collisions) which is present in the experiment but not in our PIC model. Decreases in the relative height of subharmonic spectral peaks during gas puffing observations (see Section 6.4.2 of Ref. 42) are consistent with this hypothesis.



FIG. 6. Detail from the particle simulation, showing the spectra for all three components of the electric field. Each wave shows a unique signature with respect to the different components, consistent with a parametric decay scenario. (a) The E_x component is dominated by the propagating IBW daughter wave, which is nearly electrostatic. It is a slow wave with large amplitude relative to its power content. The zero frequency spike shows that a strong DC field is also generated. (b) The E_y component is dominated by the drive frequency's fast wave; note that the amplitude of E_x and E_y at the drive frequency is the same magnitude. (c) The E_z component is very weak and shows only the ICW.

VI. NONLINEAR FLUID SIMULATIONS

Nonlinear simulations have also performed using the fluid models described in Sec. IV. We have run simulations in which the electric field amplitude of the wave is progressively increased; such simulations should manifest nonlinear effects when amplitudes become large enough. Figure 7 examines the normalized electric field spectra for two cases with wavenumber k = 220/m, for which the initial amplitudes of the electric field differ by four orders of magnitude and the initial conditions correspond to those of the linear IBW. When the initial amplitude is low (blue curve), the power spectrum is peaked only at the frequency corresponding to the linear IBW eigenmode. However, the simulation in which the initial amplitude is large has additional frequency components in the normalized power spectrum. Many of these are harmonics of the fundamental linear IBW frequency ($\omega = p\omega_0$ for integer p), consistent with the frequency doubling, and subsequent harmonic cascade phenomena one expects in the nonlinear regime. In addition to the harmonic cascade associated with the drive frequency, another IBW whose wavelength is half as large (with k = 440/m and its own fundamental frequency ω_1) is also excited. Linear combinations of these two fundamental frequencies beat together to form the rich spectrum shown in Figure 7, demonstrating that the fluid model can capture the nonlinear effects which give rise to wavenumber/frequency doubling and to beating.

In another simulation, a wave is excited in a onedimensional periodic domain by driving the x-component of the electric field with a broad spatial excitation at fixed frequency; the frequency corresponds to that of the n = 5 eigenmode in the domain (i.e., there are five wavelengths across the domain). We display the data in Figure 8; linearly, the n = 5 harmonic dominates the power spectrum. However, in the nonlinear regime, a strong zero-frequency component is



FIG. 7. Normalized power spectra for cases with low (blue) and higher (red) initial electric field amplitude in the IBW. In the low-amplitude case $[E_y(t=0) = 1 \text{ V/m}]$, the great majority of the power is focused in the fundamental IBW frequency ω_0 . In the larger-amplitude case $[E_y(t=0) = 10^4 \text{ V/m}]$, the excitation of an additional IBW with $k = 2k_0$ and its own fundamental frequency ω_1 yields a rich spectrum as the ensuing waves beat together. Multiples of the beat frequency ω_0 are also nonlinearly excited.



FIG. 8. Linear and nonlinear RF simulations which are driven at the frequency ω (drive) corresponding to the n = 5 linear eigenmode in the periodic domain. In the linear case, the n = 5 eigenmode dominates the spectrum; in the nonlinear regime, subharmonics and a strong zero-frequency component (characteristic of the PDI) arise. In fact, these spectral peaks do not arise from the PDI.

also present, together with several sub-harmonic frequencies occurring at roughly 1/5, 3/5, and 4/5 of the drive frequency. Though the presence of subharmonics and the zerofrequency peak suggest a potential PDI scenario, caution is needed; in fact, the observed subharmonics correspond precisely to allowable frequencies of n = 1, n = 3, and n = 4eigenmodes in the domain (while the strong n = 5 peak again ensues from the external drive exciting the corresponding eigenmode). While it is true that frequencies and wavelengths of these eigenmodes satisfy a PDI-like relation $[\omega(n=5) = 6.9549 \times 10^8 \text{ rad/s} = \omega(\text{drive}), \quad \omega(n=4) = 5.5784$ $\times 10^8$ rad/s, $\omega(n=1) = 1.3757 \times 10^8$ rad/s; so that $\omega(n=5)$ $\approx \omega(n=4) + \omega(n=1)$], comparison of the electric field polarizations of the linear eigenmodes suggest that these waves are unlikely to couple to one another appreciably. As well, the Ex and Ey spectra are not nearly so dissimilar as are the spectra shown in Figure 6. A more reasonable hypothesis is that nonlinear effects of the type demonstrated in Figure 7 excite the n = 1 eigenmode at its characteristic frequency $\omega(n=1)$; thereafter, sidebands about the drive frequency $[\omega(\text{drive}) \pm p \ \omega(n=1) \text{ for integer } p]$ arise via wave beating. The comparable prominence of the p = 1 sideband peaks just above and below the drive frequency lend credibility to this hypothesis, though the absence of an upshifted sideband for p = 2 is curious. In any case, the nature of the nonlinear interactions which ensue at large amplitude must be carefully scrutinized.

VII. CONCLUSIONS AND FUTURE RESEARCH

We have demonstrated that important nonlinear effects such as the parametric decay instability can be successfully simulated in the time domain, and that both fluid and PIC models can capture these effects. Such time-domain simulations afford important advantages over frequency domain modeling of complex RF propagation scenarios; they represent the only viable approach to simulate the highly nonlinear cases likely to arise when density, temperature, and current gradients are steep. Even when nonlinear interactions are weak, one need not be aware (at the problem outset in the time domain) which modes are likely to play prominent roles in the wave/plasma dynamics; instead, the dominant nonlinear interactions are revealed as the simulation proceeds. The strong similarities between PDI scenarios observed experimentally on C-Mod and our PIC simulation results, in particular, suggest that time-domain simulation methods have significant potential to contribute to our understanding of the complex nonlinear interactions at the plasma edge.

ACKNOWLEDGMENTS

We are indebted to members of the RF SciDAC project for useful discussion and feedback and for suggesting the use of J. C. Rost's C-Mod data as a suitable benchmark for the PIC approach. Dr. Rost's assistance in providing figures and data is particularly appreciated. We also acknowledge the constructive comments provided by the reviewer. This research was financially supported by the U.S. Department of Energy's SBIR program under a Phase I grant, Contract DE-SC0006242.

- ¹J. R. Myra, D. A. D'Ippolito, D. A. Russell, L. A. Berry, E. F. Jaeger, and M. D. Carter, Nucl. Fusion **46**, S455 (2006).
- ²J. Hosea, R. E. Bell, B. P. LeBlanc, C. K. Phillips, G. Taylor, E. Valeo, J. R. Wilson, E. F. Jaeger, P. M. Ryan, J. Wilgen, H. Yuh, F. Levinton, S. Sabbagh, K. Tritz, J. Parker, P. T. Bonoli, and R. Harvey, Phys. Plasmas **15**, 056104 (2008).
- ³R. L. Berger, L. Chen, P. K. Kaw, and F. W. Perkins, Phys. Fluids **20**, 1864 (1977).
- ⁴M. Porkolab, Phys. Fluids **20**, 2058 (1977).
- ⁵Y. Takase, M. Porkolab, J. J. Schuss, R. L. Watterson, C. L. Fiore, R. E. Slusher, and C. M. Surko, Phys. Fluids **28**, 983 (1985).
- ⁶V. K. Tripathi and C. S. Liu, Phys. Fluids 25, 1388 (1982).
- ⁷F. S. McDermott, G. Bekefi, K. E. Hackett, J. S. Levine, and M. Porkolab, Phys. Fluids **25**, 1488 (1982).
- ⁸F. Skiff, M. Ono, and K. L. Wong, Phys. Fluids 27, 1051 (1984).
- ⁹M. Porkolab, Fusion Eng. Des. **12**, 93 (1990).
- ¹⁰R. I. Pinsker, C. C. Petty, M. J. Mayberry, M. Porkolab, and W. W. Heidbrink, Nucl. Fusion 33, 777 (1993).
- ¹¹T. Fujii, M. Saigusa, H. Kimura, M. Ono, K. Tobita, M. Nemoto, Y. Kusama, M. Seki, S. Moriyama, T. Nishitani, H. Nakamura, H. Takeuchi, K. Annoh, S. Shinozaki, M. Terakado, and JT-60 Team, Fusion Eng. Des. **12**, 139 (1990).
- ¹²J. C. Rost, M. Porkolab, and R. L. Boivin, Phys. Plasmas 9, 1262 (2002); also available as PSFC report PSFC/JA-01-6.
- ¹³R. Van Nieuwenhove, G. Van Oost, J.-M. Noterdaeme, M. Brambilla, J. Gernhardt, and M. Porkolab, Nucl. Fusion 28, 1603 (1988).
- ¹⁴A. T. Lin and C. C. Lin, Phys. Fluids **26**, 3612 (1983).
- ¹⁵N. Xiang and J. R. Cary, Phys. Plasmas 18, 122107 (2011).

- ¹⁶J. S. Lönnroth, J. A. Heikkinen, K. M. Rantamäki, and S. J. Karttunen, Phys. Plasmas 9, 2926 (2002).
- ¹⁷E. F. Jaeger, L. A. Berry, and D. B. Batchelor, Phys. Plasmas 7, 3319 (2000).
- ¹⁸J. R. Myra, L. A. Berry, and E. F. Jaeger, in 17th Topical Conference on Radio Frequency Power in Plasmas, Clearwater, FL, May 2007.
- ¹⁹T. H. Stix, Waves in Plasmas (AIP, New York, 1992).
- ²⁰G. Chen, E. F. Jaeger, L. Berry, J. Myra, and the RF-SciDAC Team, "AORSA 1-D simulation of parametric instability (PDI) in tokamak plasmas," APS-DPP poster (2009). Abstract published in Bull. Am. Phys. Soc. 54, 71 (2009).
- ²¹A. Hakim, J. Fusion Energy 27, 36 (2008).
- ²²J. Loverich and A. Hakim, J. Fusion Energy **29**, 532 (2010).
- ²³C. Nieter and J. R. Cary, J. Comput. Phys. **196**, 448 (2004).
- ²⁴A. B. Langdon and C. K. Birdsall, Phys. Fluids 13, 2115 (1970).
- ²⁵D. L. Bruhwiler, R. E. Giacone, J. R. Cary, J. P. Verboncoeur, P. Mardahl, E. Esarey, W. P. Leemans, and B. A. Shadwick, Phys. Rev. ST Accel. Beams 4, 101302 (2001).
- ²⁶L. Sironi and A. Spitkovsky, Astrophys. J. 707, L92 (2009).
- ²⁷H. Liu, B. Wu, D. Yu, Y. Cao, and P. Duan, J. Phys. D: Appl. Phys. 43, 165202 (2010).
- ²⁸S. E. Parker, W. W. Lee, and R. A. Santoro, Phys. Rev. Lett. **71**, 2042 (1993).
- ²⁹C. C. Kim and the NIMROD team, Phys. Plasmas **15**, 072507 (2008).
- ³⁰T. M. Austin, D. N. Smithe, and V. Ranjbar, "Computing quasi-linear diffusion coefficients using the delta-f particle-in-cell method," in *Proceedings of the 18th Topical Conference on Radio Frequency Power in Plasmas*, edited by V. Bobkov and J.-M. Noterdaeme (American Institute of Physics, Melville, NY, 2009), Vol. 1187, p. 585.
- ³¹J. A. Krommes, Phys. Plasmas **14**, 090501 (2007).
- ³²A. Y. Aydemir, Phys. Plasmas 1, 822 (1994).
- ³³Z. Yu and H. Qin, Phys. Plasmas **16**, 032507 (2009).
- ³⁴D. N. Smithe, Phys. Plasmas 14, 056104 (2007).
- ³⁵K. Amyx, "GPU-Accelerated 3-D electromagnetic particle-in-cell implementations in VORPAL," in Proceedings of the GPU Technology Conference, San Jose, California, May, 2012; available at http://developer. download.nvidia.com/GTC/PDF/GTC2012/Posters/P0512_GPU_3D_EMPIC _in_VORPAL_Amyx.pdf online.
- ³⁶R. J. Procassini, C. K. Birdsall, and B. I. Cohen, Nucl. Fusion **30**, 2329 (1990).
- ³⁷E. F. Jaeger, L. A. Berry, E. D'Azevedo, D. B. Batchelor, and M. D. Carter, Phys. Plasmas 8, 1573 (2001).
- ³⁸O. Buneman, Phys. Fluids **4**, 669 (1961).
- ³⁹P. L. Colestock and R. J. Kashuba, Nucl. Fusion 23, 763 (1983).
- ⁴⁰J. C. Rost, Ph.D. dissertation, Massachusetts Institute of Technology, 1998.
- ⁴¹J. C. Rost, R. L. Boivin, M. Porkolab, J. C. Reardon, and Y. Takase, AIP Conf. Proc. 403, 85 (1997).
- ⁴²J. C. Reardon, Ph.D. dissertation, Massachusetts Institute of Technology, 1999.
- ⁴³S. J. Wukitch, "Field-aligned ICRF antenna characterization and performance in Alcator C-Mod," in presented at 54th Annual Meeting of the APS Division of Plasma Physics, November, 2012; available at http:// www.psfc.mit.edu/research/alcator/pubs/APS/APS2012/Wukitch_APS2012_ invited.pdf online.
- ⁴⁴Retaining an odd number (2p-1) of velocity moments also captures the pth cyclotron resonance, provided the species gyroradius is small compared to the relevant wavelengths in the problem; if not, corrections to the resonance of order $k_{\perp}\rho_{\alpha}$ ensue. The fluid simulations in this work satisfy this small-gyroradius constraint.